

EAST Search History

| | | | | | | |
|-----|------|--|--|----|----|------------------|
| L8 | 25 | 6 and (@rlad < "20040219" or @ad < "20040219") | US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/09/14 09:19 |
| L9 | 2588 | 8and (conditions near (sub\$1query sub adj1 query)) | US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/09/14 09:19 |
| L10 | 0 | 8 and (conditions near (sub\$1query sub adj1 query)) | US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/09/14 09:19 |
| L11 | 0 | (populat\$4 near (global adj1 condition\$3)) | US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/09/14 09:19 |
| S1 | 8 | "Floyd-Warshall" and (transitive adj1 closure) | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2006/08/03 15:52 |
| S2 | 7 | S1 and query | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2006/08/03 16:10 |
| S3 | 0 | S2 and nested | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2006/08/03 15:51 |
| S4 | 2 | S2 and global | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2006/08/03 15:51 |

EAST Search History

| | | | | | | |
|----|---|-------------------------------|---|----|----|------------------|
| S5 | 4 | "5727196".pn. "5899993".pn. | US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2006/08/03 16:11 |
|----|---|-------------------------------|---|----|----|------------------|

[Google](#)

analyzing query and subqueries

[Search](#)[Advanced Search Preferences](#)[New by default](#) [details] [Manage your web history](#)[Web](#)Results 1 - 10 of about 313,000 for analyzing query and subqueries. (0.17 seconds)

Matching and compensation tests for optimizing correlated ...

... tests for optimizing correlated **subqueries** within **query** using automatic ... for use in
 optimizing the **query**; (b) **analyzing** the **query** using matching and ...
www.patentstorm.us/patents/7167853-claims.html - 30k - [Cached](#) - [Similar pages](#)

Method of optimizing a **query** having an existi **subquery** and a not ...

(a) **analyzing** the **query** to determine whether the **query** includes a WHERE clause which
 contains an "EXISTS" **subquery** or the equivalent and a "NOT EXISTS" ...
www.patentstorm.us/patents/6529896-claims.html - 22k - [Cached](#) - [Similar pages](#)

Sub-queries supported in Analysis Services MDX

More powerful MDX in SQL Server 2005 **Analysis Services** include **query support** ... MDX
 now supports **sub-queries**, allowing you to change the criteria of the ...
searchsqlserver.techtarget.com/tip/0,289483,sid87_gci1247218,00.html - 63k -
[Cached](#) - [Similar pages](#)

SQL Queries for SQL Server - Course 534

Analyzing query plans; Enhancing **query performance**; Testing queries ... **Subqueries** in
 conditions and column expressions; Creating multilevel **subqueries** ...
www.learningtree.com/courses/534.htm - 33k - [Cached](#) - [Similar pages](#)

MySQL Bugs: #4460: Query having **subqueries** (having group by clause ...)

Replace the **subquery** with the DERIVED table in the main **query** ... preliminar test I will
 ask for my colleagues help regarding the **analyze** of your **query** ...
bugs.mysql.com/bug.php?id=4460 - 10k - [Cached](#) - [Similar pages](#)

Developing SQL Queries for SQL Server: Hands-On (Learning Tree ...)

... functions to calculate ranks; Build simple and correlated **subqueries** ... Manipulating
 hierarchical data; **Analyzing query plans** and tuning queries ...
www.seminarinformation.com/event/23875 - 28k - [Cached](#) - [Similar pages](#)

Analysis Services Query Performance Top 10 Best Practices

For more information, see Configuring the **Analysis Services Query Log** ... containing
 arbitrary shapes to reduce excessive **subqueries** where possible. ...
www.microsoft.com/technet/prodtechnol/sql/bestpractice/ssasqptb.mspx - 31k -
[Cached](#) - [Similar pages](#)

Guidelines for **analyzing** where a federated **query** is evaluated

Guidelines for **analyzing** where a federated **query** is evaluated ... Does this predicate
 contain a **subquery** involving an SQL operator that is not supported by ...
publib.boulder.ibm.com/infocenter/db2luw/v9/topic/com.ibm.db2.udb.admin.doc/doc/c0005327.htm - 9k - [Cached](#) - [Similar pages](#)

[mb-devel] optimizing a database **query**

The database was vacuumed on creation anyway, so the **analyze** step is probably
 FROM moderation_all m where m.id = 2045296; **QUERY PLAN Subquery Scan** m ...
lists.musicbrainz.org/pipermail/musicbrainz-devel/2005-January/000965.html - 14k -
[Cached](#) - [Similar pages](#)

Method and apparatus for **query** and **analysis** - Patent 6804662

Thus the **subquery** expression can be understood containing an expression (including
query and **analysis**) which, when run or evaluated, returns a document set. ...

1 [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)

Try [Google Desktop](#): search your computer as easily as you search the web.

analyzing query and subqueries

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

©2007 Google - [Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

Web Images Video News Maps Gmail more ▾

Sign in

Googleanalyzing query and subqueries global conditio Search Advanced Search PreferencesThe "AND" operator is unnecessary -- we include all search terms by default [details] [New by default](#) [Manage your web history](#)

Web

Results 1 - 10 of about 99,600 for **analyzing query and subqueries global conditions**. (0.25 seconds)**System and method for data flow analysis of complex data filters ...**A system and method for analyzing the data flow of a database query. ... Global optimization of correlated subqueries and exists predicates ...
www.patentstorm.us/patents/6785668.html - 23k - [Cached](#) - [Similar pages](#)**Global Business Solutions, Inc. - DIS100-Oracle10g Discoverer for ...**Learn Oracle's easy-to-use data access tool to view, query, analyze and report ... Multiple Conditions. 8. Queries and Subqueries • Creating Subqueries ...
gb-solutions.com/index.php?option=com_content&task=view&id=139&Itemid=35 - 20k - [Cached](#) - [Similar pages](#)**[PDF] QUERY TRANSFORMATION IN A MULTI-DATABASE ENVIRONMENT USING A ...**Argument1 : global-condition. Argument2: subquery-variable-list Transformation, Query Translation and Performance. Analysis of a New Database Computer ...
portal.acm.org/ft_gateway.cfm?id=75431&type=pdf&dl=portal&id=ACM - [Similar pages](#)**[PPT] The Web of the Future: Challenges and Opportunities for Database ...**File Format: Microsoft Powerpoint - [View as HTML](#)
Elementary conditions on element/attribute names and contents ... decompose query into subqueries; choose global evaluation order of subqueries ...
www.mpi-inf.mpg.de/~weikum/gw-edbt02-talk.ppt - [Similar pages](#)**Analysis Services Query Performance Top 10 Best Practices**For more information, see Configuring the Analysis Services Query Log. ... containing arbitrary shapes to reduce excessive subqueries where possible. ...
www.microsoft.com/technet/prodtechnol/sql/bestpractice/ssasqptb.mspx - 31k - [Cached](#) - [Similar pages](#)**The SQLite Query Optimizer Overview**There is a long list of conditions that must all be met in order for query flattening to occur. The subquery and the outer query do not both use aggregates. ...
www.sqlite.org/optoverview.html - 21k - [Cached](#) - [Similar pages](#)**[PDF] Software Agent-Oriented Frameworks for Global Query Processing**the decomposition of a query into sub-queries, each of which is sent to Query refinement. The global query analysis has to consider the semantical dis- ...
www.springerlink.com/index/X53617L785U83T71.pdf - [Similar pages](#)**GLOBAL DATABASE MANAGEMENT SYSTEM INTEGRATING HETEROGENEOUS DATA ...**The system has a complex query analysis system that is arranged to find and link For sub-queries embedded in role conditions, this is not the case. ...
www.freepatentsonline.com/EP1374095.html - 48k - [Cached](#) - [Similar pages](#)**Global database management system integrating heterogeneous data ...**[0082] In the top-level query, the user specifies the attributes for which values should be returned. For sub-queries embedded in role conditions, ...
www.freepatentsonline.com/20040083223.html - 46k - [Cached](#) - [Similar pages](#)
[More results from www.freepatentsonline.com]**SQLSquare - SQL Server Training Courses - Course Details**

Analyzing Query Plans; Enhancing Query Performance; Testing Queries ... Subqueries in conditions and column expressions; Creating multilevel subqueries ...

1 [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [**Next**](#)

Download [Google Pack](#): free essential software for your PC

analyzing query and subqueries glot

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

©2007 Google - [Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

[Web](#) [Images](#) [Video](#) [News](#) [Maps](#) [Gmail](#) [more ▾](#)
[Sign in](#)[Google](#)

analyzing query and subqueries satisfy global

[Search](#)[Advanced Search Preferences](#)The "AND" operator is unnecessary -- we include all search terms by default. [\[details\]](#)[Web](#)Results 1 - 10 of about 40,700 for **analyzing query and subqueries satisfy global conditions.** (0.19 seconds)**[PDF] Software Agent-Oriented Frameworks for Global Query Processing**

needed to **satisfy the query**. Hence, the **sub-queries** are semantically ... **Query refinement**. The **global query analysis** has to consider the semantical dis- ...
www.springerlink.com/index/X53617L785U83T71.pdf - [Similar pages](#)

The SQLite Query Optimizer Overview

There is a long list of **conditions** that must all be met in order for **query flattening** to occur.
The subquery and the outer **query** do not both use aggregates. ...
www.sqlite.org/optoverview.html - 21k - [Cached](#) - [Similar pages](#)

[PDF] QUERY TRANSFORMATION IN A MULTI-DATABASE ENVIRONMENT USING A ...

WHERE site-condition. 50. After a **subquery** has been formulated the **global query** has to be modified. In this phase, tuple variables ...
portal.acm.org/ft_gateway.cfm?id=75431&type=pdf&dl=portal&id=ACM - [Similar pages](#)

Analysis Services Query Performance Top 10 Best Practices

For more information, see Configuring the **Analysis Services Query Log**. ... containing arbitrary shapes to reduce excessive **subqueries** where possible. ...
www.microsoft.com/technet/prodtechnol/sql/bestpractice/ssasqptb.mspx - 31k - [Cached](#) - [Similar pages](#)

GLOBAL DATABASE MANAGEMENT SYSTEM INTEGRATING HETEROGENEOUS DATA ...

The system has a complex **query analysis** system that is arranged to find and are instances of a given "concept" which also **satisfy** the given **conditions**. ...
www.freepatentsonline.com/EP1374095.html - 48k - [Cached](#) - [Similar pages](#)

Global database management system integrating heterogeneous data ...

[0082] In the top-level **query**, the user specifies the attributes for which values should be returned. For **sub-queries** embedded in role **conditions**, ...
www.freepatentsonline.com/20040083223.html - 46k - [Cached](#) - [Similar pages](#)
[\[More results from www.freepatentsonline.com \]](#)

[Paper] Global Query Optimization in a Dynamic Multidatabase ...

The first alternative is to decompose a **global query** into the smallest possible **subqueries** each of which is executed by one component database system. ...
www.actapress.com/PDFViewer.aspx?paperId=14742 - [Similar pages](#)

[PDF] Towards an Exhaustive Set of Rewriting Rules for XQuery ...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Analyzing BizQuery experience and works on rewriting optimization for tra- document we can express the **query** using **subqueries** nested in XML element ...
www.ispras.ru/~grinev/mypapers/rewriting-extended.pdf - [Similar pages](#)

Analyzing, optimizing and rewriting queries using matching and ...

(a) analyzing the **query** using math and compensation between the **query** and one or more ... wherein the match function must **satisfy** the following **conditions**: ...
www.patentstorm.us/patents/6847962-claims.html - 74k - [Cached](#) - [Similar pages](#)

[PDF] Pushing quality of service information and requirements into ...

the traditional database optimizers cannot **satisfy** such The **global query** decomposition phase. generates a set of **subqueries** with location information. ...
ieeexplore.ieee.org/iel5/8622/27319/01214924.pdf?arnumber=1214924 - [Similar pages](#)

1 2 3 4 5 6 7 8 9 10 **Next**

Download [Google Pack](#): free essential software for your PC

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

©2007 Google - [Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

[Web](#) [Images](#) [Video](#) [News](#) [Maps](#) [Gmail](#) [more ▾](#)

[Sign in](#)

Google

connecting conditions analyzing queries

[Search](#)

[Advanced Search](#)

[Preferences](#)

New! [View and manage your web history](#)

[Web](#) [Books](#)

Results 1 - 10 of about 1,650,000 for **connecting conditions analyzing queries** (0.27 seconds)

OracleAS Discoverer 10g: Create Queries And Reports

Business Intelligence Tools; Discoverer Components; **Connecting to Discoverer Plus ...**

Creating **Analysis Calculations**; Using Calculations in a Condition ...

[education.oracle.com/.../show_desc.redirect?dc=D17013GC10&](#)

[p_org_id=1001&lang=US&source_call=](#) - 10k - Cached - Similar pages

Course Description

OracleAS Discoverer 10g: Create **Queries** and Reports. What you will learn: ... Creating **Analysis Calculations** Using Calculations in a Condition ...

[education.oracle.com/web_prod-plq-dad/plsql/cdesc?](#)
[dc=D17013GC10&p_org_id=1001&lang=US&p_preview=N](#) - 23k - Cached - Similar pages

[doc] OracleAS Discoverer 10g: Create Queries and Reports

File Format: Microsoft Word - [View as HTML](#)

Filter data by using **conditions** and create parameterized reports. - Use and customize

Discoverer Portlets in OracleAS Portal. - Open a workbook and **analyze** ...

[www.oracle.com/global/gr/education/cdesc/10gDCEU.doc](#) - Similar pages

d17013.html

D17013, OracleAS Discoverer 10g: Create **Queries** and Reports. Hind (ilma km-ta), 9 000.-
... Creating **Analysis Calculations** Using Calculations in a Condition ...

[www.oracle.com/global/ee/education/descriptions/ServerTools/d17013.html](#) - 66k -
Cached - Similar pages

D17013GC10

OracleAS Discoverer 10g: Create **Queries** and Reports ... Creating **Analysis Calculations**

Using Calculations in a Condition ...

[education.oracle.co.uk/html/oracle/18TR/D17013GC10.htm](#) - 8k - Cached - Similar pages

VTEC - OracleAS Discoverer 10g: Create Queries and Reports Course ...

Connecting to Discoverer Plus Data Warehouse Versus Transactional Systems ... Creating **Analysis Calculations** Using Calculations in a Condition ...

[www.vtec.org/schedule/courses/cob.asp?coID=357](#) - 16k - Cached - Similar pages

Knowledge United :: Oracle Training :: OracleAS Discoverer 10g ...

OracleAS Discoverer 10g: Create **Queries** and Reports. Course Length:, 2 days ... Creating **Analysis Calculations** Using Calculations in a Condition ...

[knowledgeunited.com/oracle/course.php?course=40](#) - 58k - Cached - Similar pages

Index: O

operators: **connecting conditions**, **Connecting conditions** with logical operators ...

optimization of **queries**: rewriting subqueries as EXISTS predicates ...

[www.i anywhere.com/developer/product_manuals/](#)

[sqlanywhere/1000/en/html/dbugen10/OIndex.html](#) - 33k - Cached - Similar pages

Qualitative Analysis for Social Scientists - Google Books Result

by Anselm L. Strauss - 1987 - Science - 336 pages

The presenter patiently answered all **queries**, while the instructor grew ... Case 5

Connecting macroscopic conditions and microscopic data As noted earlier, ...

[books.google.com/books?isbn=0521338069...](#)

Oracle10g AS Discoverer: End Users - Oracle University NZ

They will learn how to present data graphically and how to format **query** results to ...

Creating **Analysis Calculations** Using Calculations in a Condition ...

www.asparona.com/Education/DB_10gASCreateQueriesReports.aspx?CourseID=2.7 - 26k
- [Cached](#) - [Similar pages](#)

1 [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)

Try [Google Desktop](#): search your computer as easily as you search the web.

| [connecting conditions analyzing que](#)

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

©2007 Google - [Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

[Web](#) [Images](#) [Video](#) [News](#) [Maps](#) [Gmail](#) [more ▾](#)

[Sign in](#)

Google

connecting conditions analyzing queries globa

[Search](#)

[Advanced Search](#)
[Preferences](#)

New! [View and manage your web history](#)

Web Results 1 - 10 of about 1,300,000 for **connecting conditions analyzing queries global condition set** (0.27 seconds)

[DOC] OracleAS Discoverer 10g: Create Queries and Reports

File Format: Microsoft Word - [View as HTML](#)

Filter data by using **conditions** and create parameterized reports ... 5 - Setting **Query**

Management Options. 4 - 1 - Result Set Display Options ...

www.oracle.com/global/gr/education/cdesc/10gDCEU.doc - Similar pages

d17013.html

D17013, OracleAS Discoverer 10g: Create **Queries** and Reports. Hind (ilmä km-ta), 9 000.-

... Creating **Analysis** Calculations Using Calculations in a **Condition** ...

www.oracle.com/global/ee/education/descriptions/ServerTools/d17013.html - 66k -

Cached - Similar pages

Trends in Genetics : Global synthetic-lethality analysis and yeast ...

In this scenario, the 5000×5000 synthetic-lethality matrix identified using a **set** of standard growth **conditions** might be viewed as the 'base layer' of an ...

linkinghub.elsevier.com/retrieve/pii/S016895250500332X - Similar pages

[PDF] LNCS 2736 - Policy Based Enterprise (Active) Information Integration

A user poses the **queries** on the **global** schema to retrieve data, not to, may utilize a **set** of **conditions** for **analyzing** data, which were otherwise ...

www.springerlink.com/index/AX2198J0V05JMMBM.pdf - Similar pages

[PDF] Software Agent-Oriented Frameworks for Global Query Processing

Query refinement. The **global query analysis** has to consider the ... in specifying the **condition** constitute the **condition** attributes **set C**. For our **query**, ...

www.springerlink.com/index/X53617L785U83T71.pdf - Similar pages

Do Global Attractors Depend on Boundary Conditions? - Fiedler ...

The **global** attractor class is the **set** of all equivalence classes of **global** attractors ...

Adaptations to mildly nonlinear boundary **conditions** seem feasible. ...

citeseer.ist.psu.edu/fiedler96do.html - 26k - Cached - Similar pages

[PDF] Global analysis of gene function in yeast by quantitative ...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

We present a method for the **global analysis** of the function of genes in budding yeast

computed on each **set** of biological **conditions** separately and ...

www.brownlab.info/mbrown/pdf/Brown2006.pdf - Similar pages

Method and apparatus for replicating and analyzing worm programs ...

The details of the results and other activity provided by the service provider host computer

203 depend on the initial emulation **conditions** **set up** by the ...

www.freepatentsonline.com/6981279.html - 61k - Cached - Similar pages

Global Study Magazine - Site Information - Advice for students ...

Welcome to **Global** Study Magazine's terms and **conditions**, which apply to your ... in connection with the matters **set out** in these **Conditions** or otherwise. ...

www.globalstudymagazine.com/site/docs/31/ - 26k - Cached - Similar pages

[PDF] LNCS 3291 - A Necessary Condition for Semantic Interoperability in ...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

derive a necessary **condition** to foster semantic interoperability in the large and. present a method for evaluating the propagation of a **query** issued locally ...

lsirpeople.epfl.ch/aberer/PAPERS/ODBASE2004.pdf - Similar pages

1 2 3 4 5 6 7 8 9 10 **Next**

Download [Google Pack](#): free essential software for your PC

| connecting conditions analyzing que

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

©2007 Google - [Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: The ACM Digital Library The Guide

analyzing query global conditions sub queries

SEARCH

THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used: **analyzing query global conditions sub queries**

Found 66,951 of 210,707

Sort results by relevance [Save results to a Binder](#)
 Display results in expanded form [Search Tips](#)
 [Open results in a new window](#)

[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

Result page: **1** [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 Collision detection and proximity queries

Sunil Hadap, Dave Eberle, Pascal Volino, Ming C. Lin, Stephane Redon, Christer Ericson
 August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

Publisher: ACM PressFull text available: [pdf\(11.22 MB\)](#) Additional Information: [full citation](#),

methods discussed will be tied to familiar applications such as rigid body and cloth simulation, and will be compared. The course is a good overview for those developing applications in physically based modeling, VR, haptics, and robotics.

2 Exploiting perception in high-fidelity virtual environments: Exploiting perception in high-fidelity virtual environments

Additional presentations from the 24th course are available on the citation page

Mashhuda Glencross, Alan G. Chalmers, Ming C. Lin, Miguel A. Otaduy, Diego Gutierrez
 July 2006 **ACM SIGGRAPH 2006 Courses SIGGRAPH '06**

Publisher: ACM PressFull text available: [pdf\(5.07 MB\)](#) Additional Information: [full citation](#), [appendices and supplements](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

The objective of this course is to provide an introduction to the issues that must be considered when building high-fidelity 3D engaging shared virtual environments. The principles of human perception guide important development of algorithms and techniques in collaboration, graphical, auditory, and haptic rendering. We aim to show how human perception is exploited to achieve realism in high fidelity environments within the constraints of available finite computational resources. In this course w ...

Keywords: collaborative environments, haptics, high-fidelity rendering, human-computer interaction, multi-user, networked applications, perception, virtual reality

3 Special issue: AI in engineering

D. Sriram, R. Joobani
 April 1985 **ACM SIGART Bulletin**, Issue 92

Publisher: ACM PressFull text available: [pdf\(8.79 MB\)](#) Additional Information: [full citation](#), [abstract](#)

The papers in this special issue were compiled from responses to the announcement in the July 1984 issue of the SIGART newsletter and notices posted over the ARPAnet. The

interest being shown in this area is reflected in the sixty papers received from over six countries. About half the papers were received over the computer network.

4 GPGPU: general purpose computation on graphics hardware

 David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn

August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

Publisher: ACM Press

Full text available:  pdf(63.03 MB) Additional Information: [full citation](#), [abstract](#), [citations](#)

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

5 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research CASCON '97**

Publisher: IBM Press

Full text available:  pdf(4.21 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

6 A framework to support multiple query optimization for complex mining tasks

 Ruoming Jin, Kaushik Sinha, Gagan Agrawal

August 2005 **Proceedings of the 6th international workshop on Multimedia data mining: mining integrated media and complex data MDM '05**

Publisher: ACM Press

Full text available:  pdf(464.28 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

With an increasing use of data mining tools and techniques, we envision that a Knowledge Discovery and Data Mining System (KDDMS) will have to support and optimize for the following scenarios: 1) *Sequence of Queries*: A user may analyze one or more datasets by issuing a sequence of related complex mining queries, and 2) *Multiple Simultaneous Queries*: Several users may be analyzing a set of datasets concurrently, and may issue related complex queries. This paper presents a systematic ...

7 Distributed query evaluation on semistructured data

 Dan Suciu

March 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 1

Publisher: ACM Press

Full text available:  pdf(689.88 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Semistructured data is modeled as a rooted, labeled graph. The simplest kinds of queries on such data are those which traverse paths described by regular path expressions. More complex queries combine several regular path expressions, with complex data restructuring, and with sub-queries. This article addresses the problem of efficient query evaluation on distributed, semistructured databases. In our setting, the nodes of the database are distributed over a fixed number of sites, and the ...

Keywords: Distributed evaluation, nested queries, parallel complexity, regular expressions, semistructured data

8 Level set and PDE methods for computer graphics

 David Breen, Ron Fedkiw, Ken Museth, Stanley Osher, Guillermo Sapiro, Ross Whitaker
August 2004 ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04

Publisher: ACM Press

Full text available:  pdf(17.07 MB) Additional Information: [full citation](#), [abstract](#), [citations](#)

Level set methods, an important class of partial differential equation (PDE) methods, define dynamic surfaces implicitly as the level set (iso-surface) of a sampled, evolving nD function. The course begins with preparatory material that introduces the concept of using partial differential equations to solve problems in computer graphics, geometric modeling and computer vision. This will include the structure and behavior of several different types of differential equations, e.g. the level set eq ...

9 The relational model for database management: version 2

E. F. Codd
January 1990 Book

Publisher: Addison-Wesley Longman Publishing Co., Inc.

Full text available:  pdf(28.61 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#), [review](#)

From the Preface (See Front Matter for full Preface)

An important adjunct to precision is a sound theoretical foundation. The relational model is solidly based on two parts of mathematics: firstorder predicate logic and the theory of relations. This book, however, does not dwell on the theoretical foundations, but rather on all the features of the relational model that I now perceive as important for database users, and therefore for DBMS vendors. My perceptions result from 20 y ...

10 Research sessions: new styles of XML: Lazy query evaluation for Active XML

 Serge Abiteboul, Omar Benjelloun, Bogdan Cautis, Ioana Manolescu, Tova Milo, Nicoleta Preda
June 2004 Proceedings of the 2004 ACM SIGMOD international conference on Management of data SIGMOD '04

Publisher: ACM Press

Full text available:  pdf(282.38 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

In this paper, we study query evaluation on Active XML documents (AXML for short), a new generation of XML documents that has recently gained popularity. AXML documents are XML documents whose content is given partly extensionally, by explicit data elements, and partly intensionally, by embedded calls to Web services, which can be invoked to generate data. A major challenge in the efficient evaluation of queries over such documents is to detect which calls may bring data that is relevant for the ...

11 A piggyback method to collect statistics for query optimization in database management systems

Qiang Zhu, Brian Dunkel, Nandit Soparkar, Suyun Chen, Berni Schiefer, Tony Lai
November 1998 Proceedings of the 1998 conference of the Centre for Advanced Studies on Collaborative research CASCON '98

Publisher: IBM Press

Full text available:  pdf(328.82 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A database management system (DBMS) performs query optimization based on statistical information about data in the underlying data-base. Out-of-date statistics may lead to inefficient query processing in the system. Existing solutions to this problem have some drawbacks such as heavy administrative burden, high system load, and tardy updates. To overcome these drawbacks, our new approach, called the piggyback method, is proposed in this paper. The key idea is to piggyback some additional retriev ...

Keywords: access method, cost estimation, database management system, piggyback analysis, query optimization, statistics collection

12 Research sessions: query optimization: Robust query processing through

 **progressive optimization**

Volker Markl, Vijayshankar Raman, David Simmen, Guy Lohman, Hamid Pirahesh, Miso Cilimdzic

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data SIGMOD '04**

Publisher: ACM Press

Full text available:  [pdf\(331.15 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Virtually every commercial query optimizer chooses the best plan for a query using a cost model that relies heavily on accurate cardinality estimation. Cardinality estimation errors can occur due to the use of inaccurate statistics, invalid assumptions about attribute independence, parameter markers, and so on. Cardinality estimation errors may cause the optimizer to choose a sub-optimal plan. We present an approach to query processing that is extremely robust because it is able to detect and re ...

13 Research sessions: Research 10: New applications: Querying business processes

Catriel Beeri, Anat Eyal, Simon Kamenkovich, Tova Milo

September 2006 **Proceedings of the 32nd international conference on Very large data bases VLDB '06**

Publisher: VLDB Endowment

Full text available:  [pdf\(598.19 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present in this paper BP-QL, a novel query language for querying business processes. The BP-QL language is based on an intuitive model of business processes, an abstraction of the emerging BPEL (Business Process Execution Language) standard. It allows users to query business processes visually, in a manner very analogous to how such processes are typically specified, and can be employed in a distributed setting, where process components may be provided by distinct providers(peers).We describe ...

14 IS '97: model curriculum and guidelines for undergraduate degree programs in

 **information systems**

Gordon B. Davis, John T. Gorgone, J. Daniel Couger, David L. Feinstein, Herbert E. Longenecker

December 1996 **ACM SIGMIS Database , Guidelines for undergraduate degree programs on Model curriculum and guidelines for undergraduate degree programs in information systems IS '97**, Volume 28 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(7.24 MB\)](#) Additional Information: [full citation](#), [cited by](#)

15 Research session: streams and stream-based processing: Sketching streams

through the net: distributed approximate query tracking

Graham Cormode, Minos Garofalakis

August 2005 **Proceedings of the 31st international conference on Very large data bases VLDB '05**

Publisher: VLDB Endowment

Full text available:  [pdf\(273.75 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Emerging large-scale monitoring applications require continuous tracking of complex data-analysis queries over collections of physically-distributed streams. Effective solutions have to be simultaneously space/time efficient (at each remote monitor site), communication efficient (across the underlying communication network), and provide continuous, guaranteed-quality approximate query answers. In this paper, we propose

novel algorithmic solutions for the problem of continuously tracking a broad ...

16 Special issue in parallelism in database systems: Parallel query processing with zigzag trees

Mikal Ziane, Mohamed Zaït, Pascale Borla-Salamet

July 1993 **The VLDB Journal — The International Journal on Very Large Data Bases,**

Volume 2 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(1.39 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

In this article, we describe our approach to the compile-time optimization and parallelization of queries for execution in DBS3 or EDS. DBS3 is a shared-memory parallel database system, while the EDS system has a distributed-memory architecture. Because DBS3 implements a parallel dataflow execution model, this approach applies to both architectures. Using randomized search strategies enables the exploration of a search space large enough to include zigzag trees, which are intermediate between le ...

Keywords: cost function, fragmentation, pipeline, search space

17 Saturn: A scalable framework for error detection using Boolean satisfiability

 Yichen Xie, Alex Aiken

May 2007 **ACM Transactions on Programming Languages and Systems (TOPLAS),**

Volume 29 Issue 3

Publisher: ACM Press

Full text available:  pdf(742.22 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This article presents Saturn, a general framework for building precise and scalable static error detection systems. Saturn exploits recent advances in Boolean satisfiability (SAT) solvers and is path sensitive, precise down to the bit level, and models pointers and heap data. Our approach is also highly scalable, which we achieve using two techniques. First, for each program function, several optimizations compress the size of the Boolean formulas that model the control flow and data flow and ...

Keywords: Boolean satisfiability, Program analysis, error detection

18 Optimizing multiple dimensional queries simultaneously in multidimensional databases

Weifa Liang, Maria E. Orlowska, Jeffrey X. Yu

February 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 8 Issue 3-4

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(269.57 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Some significant progress related to multidimensional data analysis has been achieved in the past few years, including the design of fast algorithms for computing datacubes, selecting some precomputed group-bys to materialize, and designing efficient storage structures for multidimensional data. However, little work has been carried out on multidimensional query optimization issues. Particularly the response time (or evaluation cost) for answering several related dimensional queries simultaneous ...

Keywords: Data warehousing, MDDBs, Multiple dimensional query optimization, OLAP, Query modeling

19 Cryptography and data security

Dorothy Elizabeth Robling Denning

January 1982 Book

Publisher: Addison-Wesley Longman Publishing Co., Inc.

Full text available:  pdf(19.47 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

From the Preface (See Front Matter for full Preface)

Electronic computers have evolved from exiguous experimental enterprises in the 1940s to prolific practical data processing systems in the 1980s. As we have come to rely on these systems to process and store data, we have also come to wonder about their ability to protect valuable data.

Data security is the science and study of methods of protecting data in computer and communication systems from unauthorized disclosure ...

20 Query evaluation techniques for large databases



Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Publisher: ACM Press

Full text available:  pdf(9.37 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

 **PORTAL**
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login
 Search: The ACM Digital Library The Guide
 analyzing query global conditions sub queries satisfy conditions

THE ACM DIGITAL LIBRARY

 Feedback Report a problem Satisfaction survey

Terms used:

analyzing query global conditions sub queries satisfy conditions

Found 86,705 of 210,707

Sort results by relevance

 Save results to a Binder

Try an Advanced Search

Display results expanded form

 Search Tips

Try this search in The ACM Guide

 Open results in a new window

Results 1 - 20 of 200

Result page: **1** [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale 

- 1** [Exploiting perception in high-fidelity virtual environments: Exploiting perception in high-fidelity virtual environments](#)

 Additional presentations from the 24th course are available on the citation page

Mashhuda Glencross, Alan G. Chalmers, Ming C. Lin, Miguel A. Otaduy, Diego Gutierrez
 July 2006 **ACM SIGGRAPH 2006 Courses SIGGRAPH '06**

Publisher: ACM Press

Full text available:  pdf(5.07 MB) Additional Information: full citation, appendices and supplements, abstract, references, cited by, index terms
 mov(68:6 MIN)

The objective of this course is to provide an introduction to the issues that must be considered when building high-fidelity 3D engaging shared virtual environments. The principles of human perception guide important development of algorithms and techniques in collaboration, graphical, auditory, and haptic rendering. We aim to show how human perception is exploited to achieve realism in high fidelity environments within the constraints of available finite computational resources. In this course w ...

Keywords: collaborative environments, haptics, high-fidelity rendering, human-computer interaction, multi-user, networked applications, perception, virtual reality

- 2** [Collision detection and proximity queries](#)

 Sunil Hadap, Dave Eberle, Pascal Volino, Ming C. Lin, Stephane Redon, Christer Ericson
 August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

Publisher: ACM Press

Full text available:  pdf(11.22 MB) Additional Information: full citation, abstract

This course will primarily cover widely accepted and proved methodologies in collision detection. In addition more advanced or recent topics such as continuous collision detection, ADFs, and using graphics hardware will be introduced. When appropriate the methods discussed will be tied to familiar applications such as rigid body and cloth simulation, and will be compared. The course is a good overview for those developing applications in physically based modeling, VR, haptics, and robotics.

- 3** [Special issue: AI in engineering](#)

 D. Sriram, R. Joobbani
 April 1985 **ACM SIGART Bulletin**, Issue 92

Publisher: ACM Press

Full text available:  pdf(8.79 MB) Additional Information: full citation, abstract

The papers in this special issue were compiled from responses to the announcement in

the July 1984 issue of the SIGART newsletter and notices posted over the ARPAnet. The interest being shown in this area is reflected in the sixty papers received from over six countries. About half the papers were received over the computer network.

4 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research CASCON '97**

Publisher: IBM Press

Full text available:  pdf(4.21 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

5 Level set and PDE methods for computer graphics

 David Breen, Ron Fedkiw, Ken Museth, Stanley Osher, Guillermo Sapiro, Ross Whitaker August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

Publisher: ACM Press

Full text available:  pdf(17.07 MB) Additional Information: [full citation](#), [abstract](#), [citations](#)

Level set methods, an important class of partial differential equation (PDE) methods, define dynamic surfaces implicitly as the level set (iso-surface) of a sampled, evolving nD function. The course begins with preparatory material that introduces the concept of using partial differential equations to solve problems in computer graphics, geometric modeling and computer vision. This will include the structure and behavior of several different types of differential equations, e.g. the level set eq ...

6 GPGPU: general purpose computation on graphics hardware

 David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn

August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

Publisher: ACM Press

Full text available:  pdf(63.03 MB) Additional Information: [full citation](#), [abstract](#), [citations](#)

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

7 The relational model for database management: version 2

E. F. Codd

January 1990 Book

Publisher: Addison-Wesley Longman Publishing Co., Inc.

Full text available:  pdf(28.61 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#), [review](#)

From the Preface (See Front Matter for full Preface)

An important adjunct to precision is a sound theoretical foundation. The relational model is solidly based on two parts of mathematics: firstorder predicate logic and the theory of relations. This book, however, does not dwell on the theoretical foundations, but rather on all the features of the relational model that I now perceive as important for database users, and therefore for DBMS vendors. My perceptions result from 20 y ...

8 Cryptography and data security

Dorothy Elizabeth Robling Denning
January 1982 Book

Publisher: Addison-Wesley Longman Publishing Co., Inc.

Full text available:  pdf(19.47 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

From the Preface (See Front Matter for full Preface)

Electronic computers have evolved from exiguous experimental enterprises in the 1940s to prolific practical data processing systems in the 1980s. As we have come to rely on these systems to process and store data, we have also come to wonder about their ability to protect valuable data.

Data security is the science and study of methods of protecting data in computer and communication systems from unauthorized disclosure ...

9 Distributed query evaluation on semistructured data


Dan Suciu
March 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 1

Publisher: ACM Press

Full text available:  pdf(689.88 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Semistructured data is modeled as a rooted, labeled graph. The simplest kinds of queries on such data are those which traverse paths described by regular path expressions. More complex queries combine several regular path expressions, with complex data restructuring, and with sub-queries. This article addresses the problem of efficient query evaluation on distributed, semistructured databases. In our setting, the nodes of the database are distributed over a fixed number of sites, and the ...

Keywords: Distributed evaluation, nested queries, parallel complexity, regular expressions, semistructured data

10 Sensor networks: Fundamental scaling laws for energy-efficient storage and querying

Joon Ahn, Bhaskar Krishnamachari

May 2006 **Proceedings of the seventh ACM international symposium on Mobile ad hoc networking and computing MobiHoc '06**

Publisher: ACM Press

Full text available:  pdf(220.71 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We use a constrained optimization framework to derive fundamental scaling laws for both unstructured sensor networks (which use blind sequential search for querying) and structured sensor networks (which use efficient hash-based querying). We find that the scalability of a sensor network's performance depends upon whether or not the increase in energy and storage resources with more nodes is outweighed by the concomitant application-specific increase in event and query loads. Let m be the ...

Keywords: energy efficiency, modeling, performance analysis, querying, scalability, theory, wireless sensor networks

11 A framework to support multiple query optimization for complex mining tasks

Ruoming Jin, Kaushik Sinha, Gagan Agrawal

August 2005 **Proceedings of the 6th international workshop on Multimedia data mining: mining integrated media and complex data MDM '05**

Publisher: ACM Press

Full text available:  pdf(464.28 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

With an increasing use of data mining tools and techniques, we envision that a Knowledge Discovery and Data Mining System (KDDMS) will have to support and optimize for the following scenarios: 1) *Sequence of Queries*: A user may analyze one or more datasets by issuing a sequence of related complex mining queries, and 2) *Multiple Simultaneous Queries*: Several users may be analyzing a set of datasets concurrently, and may issue related complex queries. This paper presents a systematic ...

12 Research sessions: Research 10: New applications: Querying business processes

Catriel Beeri, Anat Eyal, Simon Kamenkovich, Tova Milo

September 2006 **Proceedings of the 32nd international conference on Very large data bases VLDB '06**

Publisher: VLDB Endowment

Full text available:  pdf(598.19 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present in this paper BP-QL, a novel query language for querying business processes. The BP-QL language is based on an intuitive model of business processes, an abstraction of the emerging BPEL (Business Process Execution Language) standard. It allows users to query business processes visually, in a manner very analogous to how such processes are typically specified, and can be employed in a distributed setting, where process components may be provided by distinct providers(peers). We describe ...

13 Research sessions: new styles of XML: Lazy query evaluation for Active XML

 Serge Abiteboul, Omar Benjelloun, Bogdan Cautis, Ioana Manolescu, Tova Milo, Nicoleta Preda

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data SIGMOD '04**

Publisher: ACM Press

Full text available:  pdf(282.38 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

In this paper, we study query evaluation on Active XML documents (AXML for short), a new generation of XML documents that has recently gained popularity. AXML documents are XML documents whose content is given partly extensionally, by explicit data elements, and partly intensionally, by embedded calls to Web services, which can be invoked to generate data. A major challenge in the efficient evaluation of queries over such documents is to detect which calls may bring data that is relevant for the ...

14 Optimization of constrained frequent set queries with 2-variable constraints

 Laks V. S. Lakshmanan, Raymond Ng, Jiawei Han, Alex Pang

June 1999 **ACM SIGMOD Record , Proceedings of the 1999 ACM SIGMOD international conference on Management of data SIGMOD '99**, Volume 28 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.65 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Currently, there is tremendous interest in providing ad-hoc mining capabilities in database management systems. As a first step towards this goal, in [15] we proposed an architecture for supporting constraint-based, human-centered, exploratory mining of various kinds of rules including associations, introduced the notion of constrained frequent set queries (CFQs), and developed effective pruning optimizations for CFQs with 1-variable (1-var) constraints. While 1-var constraints a ...

15 MiniCon: A scalable algorithm for answering queries using views

Rachel Pottinger, Alon Halevy

September 2001 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 10 Issue 2-3

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(212.60 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The problem of answering queries using views is to find efficient methods of answering a query using a set of previously materialized views over the database, rather than accessing the database relations. The problem has received significant attention because

of its relevance to a wide variety of data management problems, such as data integration, query optimization, and the maintenance of physical data independence. To date, the performance of proposed algorithms has received very little attention ...

Keywords: Data integration, Materialized views, Query optimization, Web and databases

16 Computing graphical queries over XML data

 Sara Comai, Ernesto Damiani, Piero Fraternali

October 2001 **ACM Transactions on Information Systems (TOIS)**, Volume 19 Issue 4

Publisher: ACM Press

Full text available:  pdf(707.80 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The rapid evolution of XML from a mere data exchange format to a universal syntax for encoding domain-specific information raises the need for new query languages specifically conceived to address the characteristics of XML. Such languages should be able not only to extract information from XML documents, but also to apply powerful transformation and restructuring operators, based on a well-defined semantics. Moreover, XML queries should be natural to write and understand, as nontechnical person ...

Keywords: Document restructuring, graphical query languages, semantics

17 Saturn: A scalable framework for error detection using Boolean satisfiability

 Yichen Xie, Alex Aiken

May 2007 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,

Volume 29 Issue 3

Publisher: ACM Press

Full text available:  pdf(742.22 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This article presents Saturn, a general framework for building precise and scalable static error detection systems. Saturn exploits recent advances in Boolean satisfiability (SAT) solvers and is path sensitive, precise down to the bit level, and models pointers and heap data. Our approach is also highly scalable, which we achieve using two techniques. First, for each program function, several optimizations compress the size of the Boolean formulas that model the control flow and data flow and ...

Keywords: Boolean satisfiability, Program analysis, error detection

18 Research session: streams and stream-based processing: Sketching streams through the net: distributed approximate query tracking

Graham Cormode, Minos Garofalakis

August 2005 **Proceedings of the 31st international conference on Very large data bases VLDB '05**

Publisher: VLDB Endowment

Full text available:  pdf(273.75 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Emerging large-scale monitoring applications require continuous tracking of complex data-analysis queries over collections of physically-distributed streams. Effective solutions have to be simultaneously space/time efficient (at each remote monitor site), communication efficient (across the underlying communication network), and provide continuous, guaranteed-quality approximate query answers. In this paper, we propose novel algorithmic solutions for the problem of continuously tracking a broad ...

19 A piggyback method to collect statistics for query optimization in database management systems

Qiang Zhu, Brian Dunkel, Nandit Soparkar, Suyun Chen, Berni Schieber, Tony Lai

November 1998 **Proceedings of the 1998 conference of the Centre for Advanced**

Studies on Collaborative research CASCON '98**Publisher:** IBM PressFull text available:  pdf(328.82 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A database management system (DBMS) performs query optimization based on statistical information about data in the underlying data-base. Out-of-date statistics may lead to inefficient query processing in the system. Existing solutions to this problem have some drawbacks such as heavy administrative burden, high system load, and tardy updates. To overcome these drawbacks, our new approach, called the piggyback method, is proposed in this paper. The key idea is to piggyback some additional retriev ...

Keywords: access method, cost estimation, database management system, piggyback analysis, query optimization, statistics collection

- 20 IS '97: model curriculum and guidelines for undergraduate degree programs in
 information systems

Gordon B. Davis, John T. Gorgone, J. Daniel Couger, David L. Feinstein, Herbert E. Longenecker

December 1996 **ACM SIGMIS Database , Guidelines for undergraduate degree programs on Model curriculum and guidelines for undergraduate degree programs in information systems IS '97**, Volume 28 Issue 1

Publisher: ACM PressFull text available:  pdf(7.24 MB) Additional Information: [full citation](#), [cited by](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

EAST Search History

| Ref # | Hits | Search Query | DBs | Default Operator | Plurals | Time Stamp |
|-------|------|---|--|------------------|---------|------------------|
| L1 | 3 | (analyz\$4 near quer\$4) and (conditions near (sub\$quer\$4)) | US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/09/14 09:16 |
| L2 | 2 | 1 and (@rlad < "20040219" or @ad < "20040219") | US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/09/14 09:19 |
| L3 | 0 | 2 and variables | US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/09/14 09:17 |
| L4 | 1922 | (analyz\$4 near quer\$4) | US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/09/14 09:17 |
| L5 | 1351 | 4 and (variable conditions) | US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/09/14 09:18 |
| L6 | 27 | 5 and (satisfying near condition) | US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/09/14 09:18 |
| L7 | 0 | 6 and (conditions near (sub\$1query sub adj1 query)) | US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB | OR | ON | 2007/09/14 09:19 |